

Description

SYSTEM, METHOD AND STORAGE MEDIUM FOR PREFETCHING VIA MEMORY BLOCK TAGS

BACKGROUND OF INVENTION

[0001] The invention relates to memory management and in particular, to performing generalized prefetching via memory block, or page, tags in a cache memory system.

[0002] In processing systems such as computers, the data to be utilized by a processor is stored in a memory (e.g., main memory, lower level memory) and control logic manages the transfer of data between the memory and the processor in response to requests issued by the processor. The data stored in the main memory generally includes both instructions to be executed by the processor and data to be operated on by the processor. For simplicity, both instructions and true data are referred to collectively herein as "data" unless the context requires otherwise. The time taken by a main memory access is relatively long in rela-

tion to the operating speeds of modern processors. To address this, a cache memory with a shorter access time is generally interposed between the main memory and the processor, and the control logic manages the storage of data retrieved from the main memory in the cache and the supply of data from the cache to the processor.

[0003] A typical cache is organized into multiple "lines", each line providing storage for a line of data from the main memory which may be many bytes in length. When the processor issues a request for data contained in a particular line in a page, or block, of memory, the control logic determines whether that line is stored in the cache. If the line is stored in cache (i.e., there is a cache hit), the data is retrieved from the cache. If the line is not stored in cache (i.e., there is a cache miss), the data must be retrieved from the main memory and the processor is stalled while this operation takes place. Since a cache access is much faster than a lower level memory access, it is clearly desirable to manage the system so as to achieve a high ratio of cache hits to cache misses.

[0004] Memory latency is becoming an increasingly important factor in computer system performance. An implication of this increasing importance is that cache faults from the

slowest on-chip cache are becoming more expensive in terms of performance. One approach to mitigating this problem is to increase the size of the cache. Increasing the size of the cache may improve performance, but cache memory is expensive in comparison to the slower, lower level memory. It is therefore important to use cache memory space as efficiently as possible.

[0005] One way to improve the efficiency of a cache memory system and to decrease memory latency time is to attempt to anticipate processor requests and retrieve lines of data from the memory in advance. This technique is known as prefetching. Prefetching can be performed by noting dynamic properties of the reference data stream such as sequential and/or strided accesses. Alternatively, prefetching can be performed on the basis of stored information. This stored information might be related to patterns of access within or between memory blocks or pages, or to hints produced by the compiler and/or programmer.

[0006] In order to assist in the performance of prefetching, an apparatus may store block-dependent information in main memory. This block-dependent information may be referred to as a block tag or tag. Block tags may be prepared and maintained by hardware and/or software for a

variety of purposes including that of aiding a processor in its decisions to prefetch appropriate data from memory. A distinct feature of this scheme is that it enables long term learning of computer behavior, unlike say, schemes that employ a data structure that is stored inside a processor core which necessarily is much smaller in capacity.

[0007] Given a performance goal in mind, for example, that of reducing the miss rate in a cache through prefetching, an important issue is to determine the nature of the statistical information that is to be extracted and stored in a tag, along with a representation for it that is compact, yet useful. In the same vein, methods for managing, interpreting tags and generating appropriate system commands are of prime interest. Another important issue is how this information is used and managed when there are multiple processors in a system.

[0008] The idea that knowledge of past accesses for a block, or page, in memory may be useful for preparing good prefetch candidates is well known in the art. See for instance, the reference entitled "Adaptive Variation of the Transfer Unit in a Storage Hierarchy" by P. A. Franaszek and B. T. Bennett, IBM Journal of Research and Development, Vol. 22, No. 4, July 1978. In addition, U.S. Patent

No. 6,535,961 describes a mechanism that detects bursts of access to a memory block together with the memory reference that started the burst (the "nominating line"). During this burst, memory access activity for the memory block is stored in a spatial footprint that is associated with the nominating cache line. These spatial footprints are kept in an "active macro block table." When a block becomes inactive, the corresponding spatial footprint is evicted and then stored in a "spatial footprint table." The information in the spatial footprint table is then used to issue prefetch commands.

[0009] U.S. Patent No. 6,678,795 discloses the use of a related idea to prepare prefetch candidates. An invention similar in spirit is described in U.S. Patent No. 6,134,643 and in an article by Y. Haifeng and K. Gerson entitled "DRAM-Page Based Prediction and Prefetching", 2000 IEEE International Conference on Computer Design: VLSI in computers and Processors Sept 17-20,2000 p. 267. The patent and article describe generating prefetches using the information stored in a "prediction table cache", a data structure that maintains for each block, the most recent "N" line accesses to it (each block comprises N lines) using $N \log_2 N$ bits per block entry. Further, an article by A.

Thomas and K. Gershon entitled "Distributed Prefetch-buffer/Cache Design for High Performance Memory Systems", 2nd IEEE Symposium on High Performance Computer Architecture HPCA 96, Feb 03-07 1996, p. 254, teaches a system to store, for each memory block, the address of up to some number (e.g., four) of blocks that have been referenced in the vicinity of the original block, and to use this information to generate prefetches.

[0010] Issues with the prior art described in the previous paragraphs have to do with the quality and amount of information that needs to be stored. A simplistic method that utilizes N bits to describe the accesses to a page may become polluted with irrelevant information. Maintaining the identity of the M most recently referenced lines may require M to be so large that it is a burden on storage (e.g., in the system page tables).

SUMMARY OF INVENTION

[0011] One aspect of the invention is a system for memory management. The system includes a tag cache in communication with one or more cache devices in a storage hierarchy. The tag cache includes tags of recently accessed memory blocks where each tag corresponds to one of the pages and each tag includes tag contents. The tag con-

tents control which memory lines of the corresponding memory block are prefetched into at least one of the cache devices. The tag contents are updated using a selected subset of processor references. The subset is referred to as filtered references. The tag contents are modified probabilistically at selected time or events.

[0012] Another aspect of the invention is a method for memory management. The method includes receiving a notification of a cache fault from a cache device. The notification includes a fault memory block and a fault memory line. The method also includes determining if a tag corresponding to the fault memory block is present in a tag cache. The tag includes a prefetch bit corresponding to memory lines within the memory block specified by the tag. In response to not locating the tag corresponding to the fault memory block in the tag cache, the method further includes: fetching the tag corresponding to the fault memory block into the tag cache, prefetching memory lines corresponding to prefetch bits in the tag that are set to a prefetch status into the cache device, and resetting each of the prefetch bits which were set to a prefetch status to a nonprefetch status with a selected probability. The prefetch bit corresponding to the fault memory line in

the tag is set to a prefetch status.

[0013] A further aspect of the invention is a computer program product for memory management. The computer program product includes a storage medium readable by a processing circuit and storing instructions for execution by the processing circuit for performing a method that includes receiving a notification of a cache fault from a cache device. The notification includes a fault memory block and a fault memory line. The method also includes determining if a tag corresponding to the fault memory block is present in a tag cache. The tag includes a prefetch bit corresponding to memory lines within the memory block specified by the tag. In response to not locating the tag corresponding to the fault memory block in the tag cache, the method further includes: fetching the tag corresponding to the fault memory block into the tag cache, prefetching memory lines corresponding to prefetch bits in the tag that are set to a prefetch status into the cache device, and resetting each of the prefetch bits which were set to a prefetch status to a nonprefetch status with a selected probability. The prefetch bit corresponding to the fault memory line in the tag is set to a prefetch status.

BRIEF DESCRIPTION OF DRAWINGS

- [0014] Referring now to the drawings wherein like elements are numbered alike in the several FIGURES:
- [0015] FIG. 1 is a block diagram of a system for prefetching via memory block tags in accordance with an exemplary embodiment of the present invention;
- [0016] FIG. 2 depicts the contents of a tag utilized in an exemplary embodiment of the present invention;
- [0017] FIG. 3 is a flow diagram of a process that may be utilized by an exemplary embodiment of the present invention for prefetches and prefetch bit updates for lines in a memory block; and
- [0018] FIG. 4 is a flow diagram of a process that may be utilized by an exemplary embodiment of the present invention for prefetches from a page with a proximate virtual address.

DETAILED DESCRIPTION

- [0019] An exemplary embodiment of the present invention provides a compact representation of information for use by a prefetch decision engine. The information is stored in a tag that includes N bits per memory block, or page, where N is the number of lines per block. In the rest of the document every reference to a memory block shall be under-

stood to refer not only to a contiguous portion of memory of fixed size but also to the standard notion of a system page. The terms memory block and page are utilized interchangeably in this document unless specified otherwise. Updates to the information (e.g., tags) are performed by a combination of filtering of references, along with a probabilistic aging technique (based on the use of a random number generator) that can be implemented in a relatively straight forward manner.

[0020] Additional information may also be stored in the tag to control the movement of data, such as statistics of references to blocks in proximate virtual addresses. Additional statistical information that may be stored in the tag includes the real address of the virtual page that follows the current virtual page, when the block corresponds to a page. This may be useful in determining prefetch candidates. Further information stored in the tag may include a list of processor and/or processes that have accessed the block so that process scheduling decisions can be made. In addition, external traffic conditions may be taken into account and input to prefetching decisions.

[0021] Another aspect of an exemplary embodiment of the present invention is a technique for utilizing the tags in a

multiprocessor system, where data can reside in any one of multiple caches. Each processor chip within the multiprocessor system can hold a potentially different tag, that is, the tags are not synchronized. As a processor sees a request from another processor for a specific line, it may use the information in its local version of the tag, which may have more current information, to send (push) additional lines to the requesting processor. As tags are flushed from the tag caches, a version of the tag stored in memory is updated.

[0022] An exemplary embodiment of the present invention includes a block tag format that may be utilized to perform reference aging. If only N bits are available for storage when a block consists of N lines, a first thought may be to set the i^{th} bit of the block tag to "1" whenever the i^{th} bit of the block tag is accessed. This ensures that every line of the block that has been accessed is registered as such, but has the drawback that excessive information accumulation will take place as time progresses, thus degrading the quality of the inferences that can be made with the aid of the information contained in the block tag. Instead, an exemplary embodiment of the present invention augments the above scheme with the idea of periodically

turning off each bit in the tag with a certain probability, "P".

[0023] When a microprocessor makes a memory reference that passes a filtering criterion, its corresponding block tag is retrieved from memory (if it is not already inside of the core) and is inserted in a stack, or tag cache, that is managed using replacement techniques known in the art. The prefetch bits in the tag can then be used to produce prefetches for the memory block associated with the tag. Other prefetches are possible depending on the additional features incorporated in the tag as described in the forthcoming paragraphs. Next, the probabilistic aging process is applied to the tag contents which consists of turning off every prefetch bit that is set with a certain probability P.

[0024] During the period of time that the block tag spends inside of the stack, filtered references accessed in that block are registered. Filtering refers to selecting a subset of processor references as relevant for prefetch decisions and access history. References to lines that meet the filtering criteria are used to update the corresponding block tag. When the block tag is evicted from the stack, or tag cache, the tag is stored back to memory. The lines for which a bit is set in a block tag constitute prefetch candidates the

next time that the block tag is fetched into the stack.

[0025] The amount of prefetching that is done by the system may be controlled by changing the turn-off probability parameter, "P". Another method for controlling the prefetching is to randomly select a fraction of the lines for which a bit is set to "1" and issue prefetches only for the selected ones. This method of selecting a fraction of the lines may be useful to control the amount of prefetching when an external traffic measurement indicates that only a limited amount of additional memory requests can be made.

[0026] FIG. 1 is a block diagram of a system for performing prefetching via memory block tags in accordance with an exemplary embodiment of the present invention. The system depicted in FIG. 1 includes two processor subsystems P_a 116 and P_b 108 along with their dedicated level one (L1) and level two (L2) caches. The dedicated L1 cache for subsystem P_a 116 is denoted as $L1_a$ cache 114 and the dedicated L2 cache is denoted as $L2_a$ cache 112. Similarly, the dedicated L1 cache for subsystem P_b 108 is denoted as $L1_b$ cache 106 and the dedicated L2 cache is denoted as $L2_b$ cache 104. For purposes of discussion, it is assumed that the data in memory 102 is partitioned into blocks, or pages, of four thousand and ninety six (4K)

bytes, that each block, or page, holds thirty-two (32) lines and that each line contains one hundred and twenty-eight (128) bytes of data. Also depicted in FIG. 1 is tag cache_a 118 for subsystem P_a 116 and tag cache_b 110 for subsystem P_b 108. The tags in tag cache_a 118 and tag cache_b 110 are entities associated with individual blocks in the memory 102 and are described below in reference to FIG. 2.

[0027] As is shown in FIG. 1, the subsystem P_a 116 is in communication with the L1_a cache 114 which in turn is in communication with the L2_a cache 112. The L2_a cache 112 is in communication with the memory 102. As is known in the art, the communication between these components may be via a direct connection or a network connection. The communication path between the subsystem P_a 116, L1_a cache 114, L2_a cache 112 and memory 102 is utilized to request and receive data. The data requested may be stored in the L1_a cache 114, the L2_a cache 112 or the memory 102. In addition, FIG. 1 depicts tag cache_a 118 in communication with both the L1_a cache 114 and the L2_a cache 112. In an exemplary embodiment of the present invention, the tag cache_a 118 includes tags of blocks that hold lines on which the L1_a cache 114 has recently

missed. The tag cache_a 118 is updated as lines within these blocks miss from the L1_a cache 114. Further, the L2_a cache 112 is in communication with the L2_b cache 104 to allow the subsystems P_a 116 and P_b 108 to share access to data in their respective L2 caches in order to avoid the need for satisfying a data request from memory 102.

[0028] Similarly, the subsystem P_b 108 is in communication with the L1_b cache 106 which in turn is in communication with the L2_b cache 104. The L2_b cache 104 is in communication with the memory 102. As is known in the art, the communication between these components may be via a direct connection or a network connection. The communication path between the subsystem P_b 108, L1_b cache 106, L2_b cache 104 and memory 102 is utilized to request and receive data. The data requested may be stored in the L1_b cache 106, the L2_b cache 104 or the memory 102. In addition, FIG. 1 depicts tag cache_b 110 in communication with both the L1_b cache 106 and the L2_b cache 104. In an exemplary embodiment of the present invention, the tag cache_b 110 includes tags of blocks that hold lines on which the L1_b cache 104 has recently missed. The tag cache_b 110 is updated as lines within these blocks miss from the L1_b cache 104. The components depicted in FIG.

1 are shown as separate devices, however as is known in the art, all or a subset of these components may be included in the same device.

[0029] Each tag in tag cache_a 118 and tag cache_b 110 is associated with an individual block, or page, in memory 102. In an exemplary embodiment of the present invention, tag cache_a 118 and tag cache_b 110 are organized as standard cache structures with the storage of tags divided into a set of equivalence classes. Searches for tags associated with a given block in lower level memory are performed in any manner known in the art for cache memory (e.g., by associative search within a congruence class). In an exemplary embodiment of the present invention, it is assumed, for simplicity, that tag cache_a 118 and tag cache_b 110 are fully associative. Any structure known in the art may be utilized to implement tag cache_a 118 and tag cache_b 110. Computer instructions to implement the processes described herein may be located on one or both of the tag caches, on a memory controller and/or on a processor. As is known in the art the computer instructions may be located on one device or distributed among multiple devices.

[0030] FIG. 2 depicts the contents of a tag 202 in an exemplary

embodiment of the present invention. Both tag cache_a 118 and tag cache_b 110 utilize the tag structure depicted in FIG. 2. A tag 202 includes a page real address field 204, a prefetch bits field 206 and a next virtual page prefetch bit field 208. The page real address field 204 refers to the address of the block, or page, in memory 102. The prefetch bits field 206 includes one bit for each line in the block, or page. As described previously, for pages that are 4K, with lines of one hundred and twenty-eight (128) bytes, there are thirty-two (32) lines per page. Therefore, there are thirty-two (32) bits included in the prefetch bits field 206, each corresponding to a different line in the page. Bits are set to "1", a prefetch status, if the line has been referenced during a current visit to the tag cache (e.g., tag cache_a 118, tag cache_b 110) and the reference has passed through the filtering process. Otherwise the bits are set to "0", a nonprefetch status. The next virtual page prefetch bit field 208 (also referred to as the next virtual memory block bit) indicates if lines contained in the next block, or page, in the virtual address space should be prefetched along with the current page specified by the tag in the page real address field 204.

[0031] An exemplary embodiment of the present invention oper-

ates as follows. Given certain references by a processor P_a 116 (the discussion will refer to subsystem P_a 116, however, the same principles apply to subsystem P_b 108) to a line, the tag entries associated with the page may be updated or referenced. If at the time of reference, the tag 202 is not present in the tag cache a 118, the tag 202 is fetched from memory 102, possibly displacing another tag 202 currently in the tag cache a 118. The displaced tag 202 is written back to memory 102 without regard to the values of tags 202 corresponding to the same page that are held by other processors. An example would be a tag 202 held by tag cache b 110 for subsystem P_b 108 that corresponds to the same page in memory 102. In an alternate exemplary embodiment of the present invention, the values of the tags 202 corresponding to the same page and held by other processors are taken into account when storing the displaced tag 202.

[0032] At the time that a tag 202 is fetched from memory 102 and inserted into the tag cache a 118, certain lines from the associated page, as described by the bits in the prefetch bits field 206 in the tag 202, may be prefetched. These lines are herein denoted as prefetchable. This process is described in reference to FIG. 3 below. In addition,

the value in the next virtual page prefetch bit field 208 is examined, and if it is set to "1", or a prefetch status, lines from the next virtual page are prefetched as described in reference to FIG. 4 below. In addition, the tag 202 may also be updated to reflect current information regarding references to pages which are proximate in the virtual address space. In an exemplary embodiment of the present invention, fetching a line from the L2 cache of another processor may cause the other processor to send not just the requested line, but also all the lines that are indicated as prefetchable in the tag held by the other processor.

[0033] FIG. 3 is a flow diagram of a method that may be utilized by an exemplary embodiment of the present invention for performing prefetches and prefetch bit updates for lines in a page. At step 302, the system is notified that a L1 cache miss to a line, in a particular page, has occurred. At step 304, a filtering process is applied to determine if any action should be taken by the tag cache_a 118 in response to the notification. A filter that selects which microprocessor (either local or external) references are used in preparation of the tags 202 may be utilized to enhance the operation of the device. For illustrative purposes, consider a microprocessor device in which prefetches from an exter-

nal memory are to be inserted in an L2 cache, and therefore, the purpose of the prefetching device is to reduce the L2 miss rate. The entire set of references by the processor to cache lines may contain redundant or unnecessary information, while insufficient information may be provided by just considering L2 misses. An exemplary embodiment of the present invention provides a mechanism that is intermediate between the two extremes. The natural geometry of a set associative cache is exploited to implement a "virtual" filter that requires no additional hardware. For example, the references that cross a prescribed threshold in the stack order of a least recently used (LRU) managed set associative cache (for example L2) are the ones that are used for the preparation of the block tag information. Here, every reference to a line that makes a transition from being the most recently used (MRU) to the second MRU in the stack causes the bit corresponding to this line in the tag to be set to "1". Lines with entries set to "1" are prefetchable.

[0034] At step 306, it is determined if the tag 202 corresponding to the line and page is present in the tag cache_a 118. If the tag 202 corresponding to the line and page is present in the tag cache_a 118, then step 308 is performed to set a

bit in the tag corresponding to the line to signify that the line was referenced. This is done by setting a bit corresponding to the line to "1" in the prefetch bits field 206. In an exemplary embodiment of the present invention, the bit in the tag corresponding to the line is set to on to signify that the line was referenced whenever a cache device accesses the line (e.g., for read, update) and not just in response to a fault occurring. In this manner, bits that correspond to lines that are prefetched into the cache are also set to "1" when the line is accessed by the cache. After step 308 is performed, processing ends at step 310. Alternatively, if the tag is not present in the tag cache_a 118, as determined at step 306, then step 312 is performed.

[0035] At step 312, the tag 202 corresponding to the page is fetched into the tag cache_a 118 from memory 102. Step 312 includes updating the next virtual page prefetch bit field 208, in the tag corresponding to a previous page. At the time that the tag is fetched into the tag cache_a 118, the processor checks the tag cache_a 118 for the presence of a tag 202 corresponding to the page which preceded the current one in its virtual address space. If the page is present, the count in the next virtual page prefetch bit

field 208 of the preceding page tag 202 is set to "1". At the time that the tag 202 is written back to memory 102, the next virtual page prefetch bit field 208 is decremented to "0", or to a nonprefetch status, with a probability P. Such decrementation uses the same logic as that for decrementing the prefetch bits field 206. In this example, the threshold for prefetching the tag 202, and thus the lines, from the neighboring pages is set at "1".

[0036] Processing then continues to both step 314 and step 316. At step 314, the next page prefetch logic as described in reference to FIG. 4 is performed. At step 316, lines from the page are prefetched from memory 202 into the L2_a cache 112 if their corresponding bit in the prefetch bits field 206 is set to "1". Next, at step 318 in FIG. 3, each prefetch bit in the prefetch bits field 206 is set to "0", or a nonprefetch status, with a probability of "P" to implement the concept of reference aging. In an exemplary embodiment of the present invention, the probability, P, is approximately one out of eight (1/8). P may be set to any percent value from zero to one and may be utilized to assist in controlling the amount of prefetch activity. This process of aging augments the reference bits in the prefetch bits field 206 in the tag 202 with a procedure

which decrements (reduces to zero) each "1" bit with some stated probability "P" each time a tag is written back to memory from the tag cache_a 118.

[0037] In the exemplary embodiment described in reference to FIG. 3, the aging process is performed when the tag 202 is fetched into the tag cache_a 118. This has the same effect as performing the aging process each time a tag is written back to memory. An exemplary embodiment of the present invention keeps track of the address which represents the most recent reference to a page different from the tag being processed. For this address, the three least significant bits are examined. If all are zero, then the decision is to decrement to "0". This is an example of one method of producing a pseudo random bit that is "1" with a certain probability. After step 318, processing continues to step 308, as discussed previously.

[0038] FIG. 4 is a flow diagram of a method for prefetches from a block, or page, with a proximate virtual address that may be utilized by an exemplary embodiment of the present invention. The processing in FIG. 4 is invoked from step 314 in FIG. 3. At step 402, a check is made to determine if the next virtual page prefetch bit field 208 (also referred to as the next virtual memory block bit) for the tag con-

tains a "1". If it does not contain a "1", signifying that the next virtual page should not be prefetched, processing ends at step 404. Alternatively, if the next virtual page prefetch bit field 208 for the tag does contain a "1", then step 406 is performed to fetch the next virtual page tag 202 into the tag cache_a 118. Next, at step 408, the indicated lines in the tag 202 are prefetched into the L2_a cache 112. As discussed previously, the lines are prefetched if they contain the value "1" in their corresponding bit in the prefetch bits field 206 in the tag 202. Next, each bit in the prefetch bits field 206 that contains a "1" is set to a "0" with a probability of P. In an exemplary embodiment of the present invention, the probability, P, is approximately one out of eight (1/8). As discussed previously, P may be set to any percent value from zero to one and may be utilized to assist in controlling the amount of prefetch activity. Although in this embodiment the probability P for modifying the line and next virtual page prefetch bit in a tag is the same, it should be understood that in some applications, these probabilities should be different. The processing then ends at step 412.

[0039] An exemplary embodiment of the present invention provides a compact representation of information for use by

a prefetch decision engine. The compact size of the tags may lead to less storage space being required and to increased performance of the cache system. In addition, the ability to age the data stored in the tag may lead to a higher cache hit ratio. An exemplary embodiment of the present invention includes the ability to control the amount of prefetching performed by changing the value of the probability parameter that is used by the aging algorithm. This can assist in balancing the amount of data that is prefetched with the amount of available cache storage.

[0040] As described above, the embodiments of the invention may be embodied in the form of computer implemented processes and apparatuses for practicing those processes. Embodiments of the invention may also be embodied in the form of computer program code containing instructions embodied in tangible media, such as floppy diskettes, CD-ROMs, hard drives, or any other computer readable storage medium, wherein, when the computer program code is loaded into and executed by a computer, the computer becomes an apparatus for practicing the invention. An embodiment of the present invention can also be embodied in the form of computer program code, for

example, whether stored in a storage medium, loaded into and/or executed by a computer, or transmitted over some transmission medium, such as over electrical wiring or cabling, through fiber optics, or via electromagnetic radiation, wherein, when the computer program code is loaded into and executed by a computer, the computer becomes an apparatus for practicing the invention. When implemented on a general-purpose microprocessor, the computer program code segments configure the microprocessor to create specific logic circuits.

[0041] While the invention has been described with reference to exemplary embodiments, it will be understood by those skilled in the art that various changes may be made and equivalents may be substituted for elements thereof without departing from the scope of the invention. In addition, many modifications may be made to adapt a particular situation or material to the teachings of the invention without departing from the essential scope thereof. Therefore, it is intended that the invention not be limited to the particular embodiment disclosed as the best mode contemplated for carrying out this invention, but that the invention will include all embodiments falling within the scope of the appended claims. Moreover, the use of the

terms first, second, etc. do not denote any order or importance, but rather the terms first, second, etc. are used to distinguish one element from another.